DØ Computing and Software Operations and Upgrade Plan

DØ Collaboration

August 28, 2003

CHAPTER 1 – Introduction

This document presents a status report on DØ Computing and Software operation as well as an updated computing and software plan. The update covers the years 2004-2006 and includes scope and cost estimates for hardware and software upgrades. For the years 2004-2005 the document details equipment spending to cover both the operation and upgrade of the existing system[1]. The document also describes our successful and increasing use of remote computing resources.

Since the June 2002 Bird review, the DØ collaboration has improved the operational stability of the computing systems and capitalized on our global computing model. We are successfully accumulating, storing, and analyzing data in support of the experimental program. The computing team has satisfied the needs of the experiment through two major analysis cycles. As a result, the DØ collaboration has a more detailed understanding of the computing needs for analysis, event reconstruction time, and the re-reconstruction (or reprocessing) needs. This knowledge has been folded into our cost estimates.

A number of projects and innovations have reached completion or are well underway. Sequential Access by Metadata (SAM) is in use at all official DØ computing sites. We are working with CDF and the Computation and Communication Fabric (CCF) Department to integrate dCache, which will provide a tapeless data path and buffering to better support offsite operations. Based in part on work with CDF on SAM, the path to LINUX-only operation for SAM is understood. DØ has begun to reduce the reliance on the SGI O2000 (DØmino) by commissioning the Central Analysis Backend (CAB). The number of processors on DØmino has been reduced to 128. Infrastructure applications such as Runjob and DØtools for user job submission are in common use. We are deploying a runtime environment. As a result of these initiatives the functionality of the reconstruction and simulation programs and trigger capabilities have increased dramatically.

The DØ Computing and Software model relies on contributions from collaborating institutions. For example, the complete event simulation chain including event generation, detector simulation, digitization, reconstruction, and triggering takes place offsite at remote production centers. The DØ Regional Analysis Center Working Group studied requirements and potential organizations to facilitate remote analysis, and the analysis was taken further by the Offsite Analysis Task Force which oversaw efforts at IN2P3 and GridKa centers to understand performing analysis remotely. The Task Force also considered the financial contributions of each collaborating institution, leading to a proposed model of contributions based on the costs of implementing a "Virtual Center" at

[1] Originally we used the laboratory's luminosity profile from Associate Director Steve Holmes' January 2002 talk to HEPAP as an input to our planning. For this update, we have based projections mainly on our experience so far in Run II. The number of events we are writing to tape is essentially independent of accelerator luminosity, though we have factored in the expected increase in the complexity of the events as the instantaneous luminosity increases.

FNAL. In addition, we are beginning a large scale reprocessing of data at the remote centers, starting with using the National Partnership for Advanced Computing Infrastructure (NPACI) resources at University of Michigan and expanding to sites in Europe. Global considerations are fundamental to the next evolution of the DØ computing model and we have changed the DØ Computing and Core Software organization to reflect this. We are in the process of developing a Grid strategy intended to meet our production and analysis needs. The strategy will accommodate technical constraints, our production needs, our existing tools and political realities.

1.1 Overview of the Computing Model and Systems

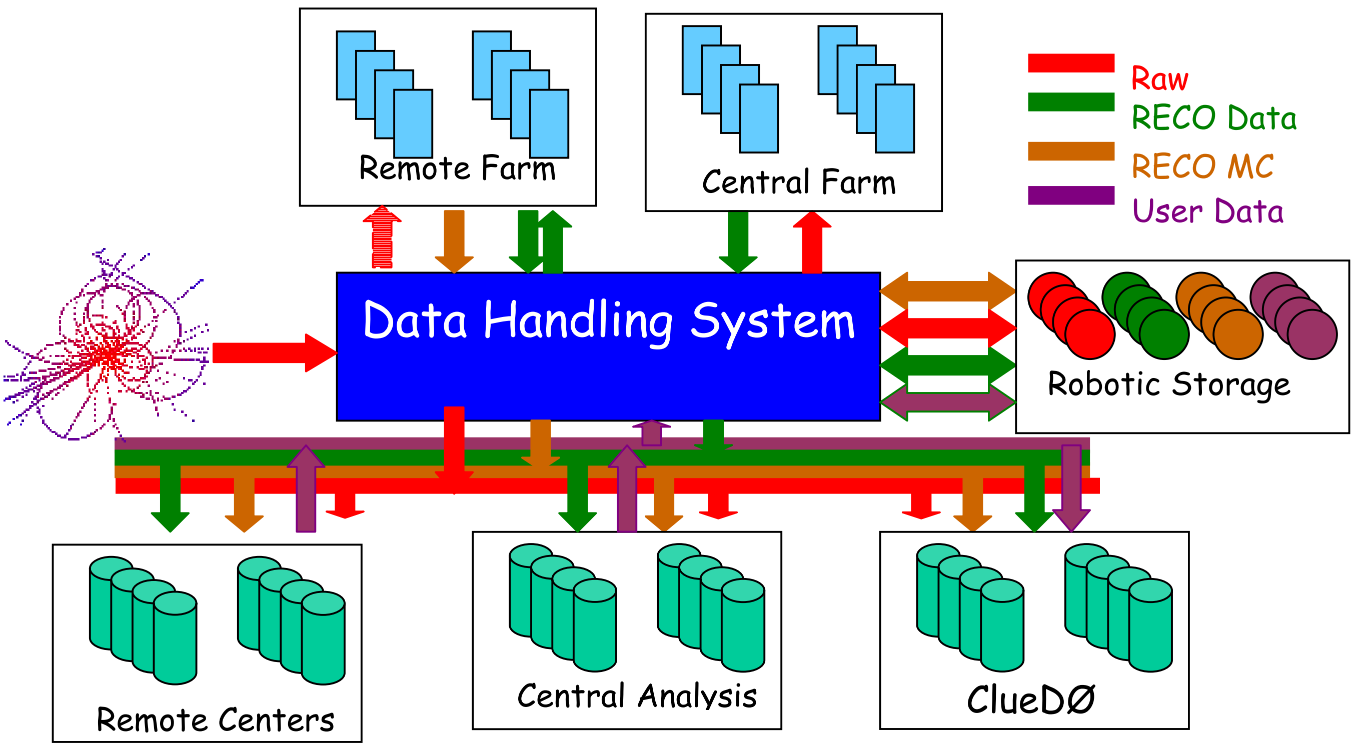

Figure 1.1 shows a schematic of the elements of the DØ computing systems, shown as linked via the data handling system, SAM. The data flow through the system is color-coded as follows: raw data is red, reconstructed data is green, Monte Carlo is brown, with user data shown in purple.

Figure 1.1 shows a schematic of the elements of the DØ computing systems. All requests for data from all systems (farm production to user requests) are handled by the Sequential Access via Metadata (SAM) system. The DØ detector/online is represented as an event display, with raw data in red, sent from the online system to the robotic storage, read from tape as raw data to the central FNAL farm, and then stored back into robotic storage from the farm as reconstructed data (shown in green). Raw data also moves from tape through the system to the user systems. In additional, all MC simulation (shown in brown) takes place at remote farms, and as stated above, re-reconstruction efforts have started. Reconstructed data or simulated data is available to remote farms for re-reconstruction and to users for analysis, and is also stored in robotic storage. After implementing the calibration database proxy servers, we can also pursue reconstruction from raw data at the remote farms. User data (shown in purple) is available at the remote analysis centers, the Central Analysis System and the DØ desktop cluster, CLuEDØ. The Central Analysis System consists of an SGI machine (DØmino) used as a file server to Linux computing nodes referred to as CAB (Central Analysis Backend). The remote centers enable people who are not resident at FNAL to have access to the data, with routing to the centers through the Central Analysis System. The remote center implementation is not directly specified—it can be a hub and spoke system with a large center serving out to smaller desktop clusters or users can directly use the large center, at the discretion of those who control the resources. There are centers in the Canada (WestGRID), France (IN2P3), Germany (GridKa at Karlsruhe), Southwestern US (centered at University of Texas at Arlington), and the United Kingdom.

1.2 Management of DØ Computing and Core Software

The DØ organization chart can be obtained at
http://DØserver1.fnal.gov/Projects/UpgradeProject/organization/Organization_DØ_public.pdf

DØ Computing and Core Software has three management branches, Global Operations, Infrastructure, and Projects. In addition, the Computing Planning Board (CPB) advises the DØ collaboration management and represents the collaboration on computing issues. The CPB makes and implements strategic decisions with respect to computing, and includes the Deputy Physics Coordinator and the Algorithms Group Deputy in the membership. The globally distributed nature of DØ (and DØ Computing) is a consideration when filling the management and CPB positions. The CPB also administers the "DØ Virtual Center", used as a cost basis for estimating financial contributions to DØ via supplying computing resources to the collaboration.

The DØ Computing and Core Software Leader is also the DØ Computing and Analysis Department Head. DØ works closely with the FNAL Computing Division on a number of joint projects, including SAM and uses CD products to good effect. CD also supplies crucial (and excellent) hardware and software support in a number of areas.

The local administration of the remote facilities is handled in different ways for different centers; however, all centers doing production work must supply a point of contact for the appropriate Global Operations Group. System support issues across the DØ community are currently handled via mail list, however, we anticipate setting up a more formal structure.

In addition to the computing administration of the remote centers, DØ collaborators have been investigating the concept of a "Regional Center", which is a self-organized sociological construct to improve and support doing analysis away from FNAL.

CHAPTER 2 –Reconstruction and Simulation

2.1 Status of Executables

In the last year, the major software components of the experiment such as reconstruction, simulation, and Level 3 trigger have undergone significant improvements.

The reconstruction has undergone considerable evolution. The program had been developed on the basis of ideal detector simulations. Completion of detector commissioning has resulted in a better understanding of the actual detector performance and capabilities. This is reflected in the inclusion in the reconstruction of alignment and calibration constants for the sub-detectors and tuning of parameterizations and uncertainties all based on data. Strategies have been developed to cope with deficiencies of subsystems. For example, algorithms have been developed to correct the effect of non-linearities in the calorimeter electronics and to reduce sensitivity to electronic noise.

A profound revision of the tracking algorithm has been implemented, resulting in major improvements in speed, efficiency, fake rate, and the ability to reconstruct low-p_T and large-impact-parameter tracks. These improvements will allow us to access samples of great interest (e.g. pions from kaon decay, electrons from photon conversion, J/ψ or Υ resonances) and make another step in the understanding of the detector. The preshower detectors have been added to the reconstruction program. Muon identification has been enhanced by combining central tracking and muon spectrometer information. Lastly, the most compact data tier, the thumbnail, has been delivered and is now the basis for analysis.

To further improve the reconstruction algorithms and to perform precision measurements, realistic detector simulation is mandatory. Significant progress has been made by the inclusion of surveyed detector geometry, a better description of the magnetic field map, the effects of electronic noise and inefficiencies based on real data, and simulation of the non-linearities of the calorimeter electronics. For a realistic model of the electronic noise and the effects of pileup and multiple interactions, the simulation is now capable of overlaying collider data events and generated physics events.

2.2 Performance Characteristics of the Offline Executables

In the 2002 plan, the projections shown in Table 2.1 for the reconstruction time per event were used. These projections were based on low luminosity data processed with version p10 of the reconstruction program. The calculations were scaled to higher luminosities using Monte Carlo estimates. Timing data for the two most recent versions of the reconstruction program, p13 and p14, are shown in Figure 2.1. The two versions have very similar performance, which is encouraging in light of the significantly improved capabilities of p14 relative to p13. However the DØ reconstruction program has become as slow as the high luminosity estimates made 18 months ago. In addition, the memory usage of the reconstruction program is such that the older dual-CPU nodes at the remote farms producing Monte Carlo are restricted to running a single process. There is active progress in understanding some of the classes of events that take a disproportionately long time to process. The current version of the reconstruction running on the production farms is p14.03, and many operational fixes will be available in p14.04, which is in preparation.

Operationally, a significant different between the p13 and p14 versions was the deployment of the calibration data base servers to extract the constants from the databases. As expected, the commissioning period did lead to some operational problems, many of which are understood and have been addressed. The performance of the calibration database servers is monitored, and can be seen at *http://dbsmon.fnal.gov/d0/d0.html*. The servers for the farms and for general usage are monitored separately.

Instantaneous Luminosity ($cm^{-2}sec^{-1}$)	Estimated Reconstruction processing time (GHz) (sec/event)
9e31	13
20e31	18
50e31	40
50e31 (132 nsec crossing[2])	16

Table 2.1 shows the 2002 estimated reconstruction time for various points of instantaneous luminosity. We used the p10 measurements for data, added the particle ID times for Monte Carlo Z events, and scaled based on the known tracking performance as a function of number of interactions.

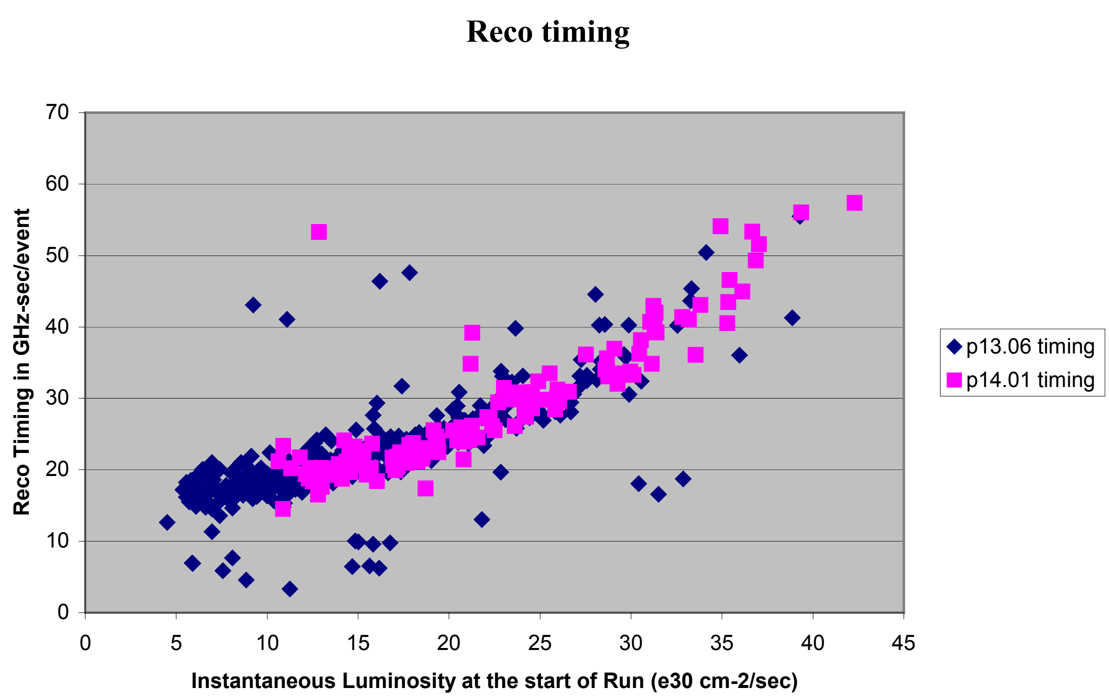

Figure 2.1 shows the measured reconstruction time as a function of instantaneous luminosity. The blue points show the time for p13 and the pink points show the time for p14. As can be seen, the timing and shape for these two versions is the same.

The output size per event for the data summary tape (DST) and the thumbnail (TMB) are both roughly double the specification from the 2002 estimates. In the case of the TMB, this is due to the fact that we have chosen to keep calorimeter cell information while the calorimeter performance is being studied. This choice has been well justified in that it enabled important corrections to be applied at the end level analysis stage. In the case of

[2] Running at 132 nanosecond crossing time is no longer within the scope of the Tevatron program.

the DST, the current size in p14 is 300 Kb/event, which is not sustainable. We assume for the planning purposes that the DST size will return to the p13 size of 200 Kb/event.

The Monte Carlo performance can be seen in the table below. These results include an average of 0.5 interactions, in addition to the hard scattering and with simulation done with plate-level GEANT. For these two representative samples about 300-350 seconds are required per event.

Process Time per Event (sec/event)	**Generation**	**Detector Simulation**	**Digitization**	**Reconstruction**	**Analyze**	**Total**
WW inclusive	0.8	280	20	19	4.5	325
Technirho	0.8	300	20	21	5	345

Table 2.2 shows current Monte Carlo chain generation time per event on a 500 MHz machine for plate level samples. Note to committee: this example is from 2002, and will be updated prior to the review—it is incomplete because we now include the trigger simulation in the full chain, and includes a p10 reconstruction time.

Parameterized MC time/event	*1 GHz-sec/event*
Full Geant Chain MC time/event	*170 GHz-sec/event*
Reconstruction on collider data time/event	*50 (60,80) GHz-sec/event*
Data DST size/event	*200 Kbytes/event*
Data TMB size/event	*25 Kbytes/event*
MC Døgstar size/event	*700 Kbytes/event*
MC Døsim size/event	*300 Kbytes/event*
MC DST size/event	*200 Kbytes/event*
MC TMB size/event	*25 Kbytes/event*

Table 2.3 shows the planning assumptions based on the simulation and reconstruction parameters. The three numbers in the Reco time/ event reflects 2004, 2005 and 2006 planning, and assumes that the experiment takes action to achieve that performance.

The information in Table 2.3 is derived from our production experience. The reconstruction times used for planning are 50 GHz-sec in 2004, 60 GHz-sec in 2005 and 80 GHz-sec in 2006. As can be seen in Figure 2.1, these numbers assume current performance in 2004, with a year to speed the code up, and with rising instantaneous luminosities in 2005 and 2006. In the event of a long shutdown in 2006, we would likely perform re-reconstruction of the entire data set, and thus still need to expand the FNAL farm.

CHAPTER 3 – Data Handling

The DØ collaboration uses the Sequential Access via Metadata (SAM) system developed jointly by the FNAL computing division and DØ. SAM provides data access through an interface that specifies the metadata appropriate to the data sample. SAM determines which files meet the criteria and then delivers that set of files to an application. A SAM application called the station communicates with a user request (called a project) and communicates with the database to determine both the file location and how to transfer those files to a user's job. Files that have been opened (or closed) by a user's application are counted as "consumed" data, and information on files consumed by a project is recorded in the database. Data can be delivered to a user application from within the station cache (either directly or via rcp transfer in a distributed cache), retrieved from another station's cache, or brought into the station cache from the ENSTORE mass storage system.

The combination of data access via meta-data, internal caching mechanisms, wide and local area transfer protocols, and storage elements leads to a uniform, collaboration wide tool underpinning all of the DØ computing. The database tracks usage as well as catalogs the events and files. Extensive plots of the SAM system performance are available online at http://DØdb.fnal.gov/sam/plots-and-stats.html. Figure 3.1 shows one such plot—the transfers in and out for all stations for a recent 30 day period.

The data handling system is monitored in a number of ways and at three levels. The collaboration conducts shifts of 16 hours duration per day to monitor the overall health of the system, answer user-level questions, diagnose problems, and restart systems when necessary. The shifters have been increasingly effective, as can be seen from Figure 3.2, which shows that the shifters are fielding an increasing fraction of e-mail queries as a function of time. The shifters are also finding and solving problems before e-mail traffic is generated, but we have not quantified this yet. In the event that the shifter cannot solve the problem, the next level of intervention is the SAM expert on call, who either fixes problems or contacts the developers, who represent the third level of monitoring. One of the prime diagnostic tools for the major stations is SAM-TV, which shows the health of a system at a glance. A snapshot of SAM-TV is shown in Figure 3.3. Other monitoring tools include the Enstore and FBSNG web pages, as well as other SAM pages.

In addition to monitoring the system on a real-time basis, DØ keeps weekly statistics to monitor the overall use of the analysis systems and to draw conclusions about system usage. The metrics are tracked for the central-analysis station (which runs on DØmino), the CAB station (which runs on the central analysis backend of linux nodes), the CLuEDØ station (which runs on the DØ desktop cluster), and the DØKarlsruhe station, (which runs at GridKa in Germany). In Figures 3.4 and 3.5, the amount of data consumed as a function of time is plotted for these four stations expressed as the number of events and Terabytes. This combination of stations monitors several types of user activity. As a note, during one week, 1.4 billion events were consumed on CAB.

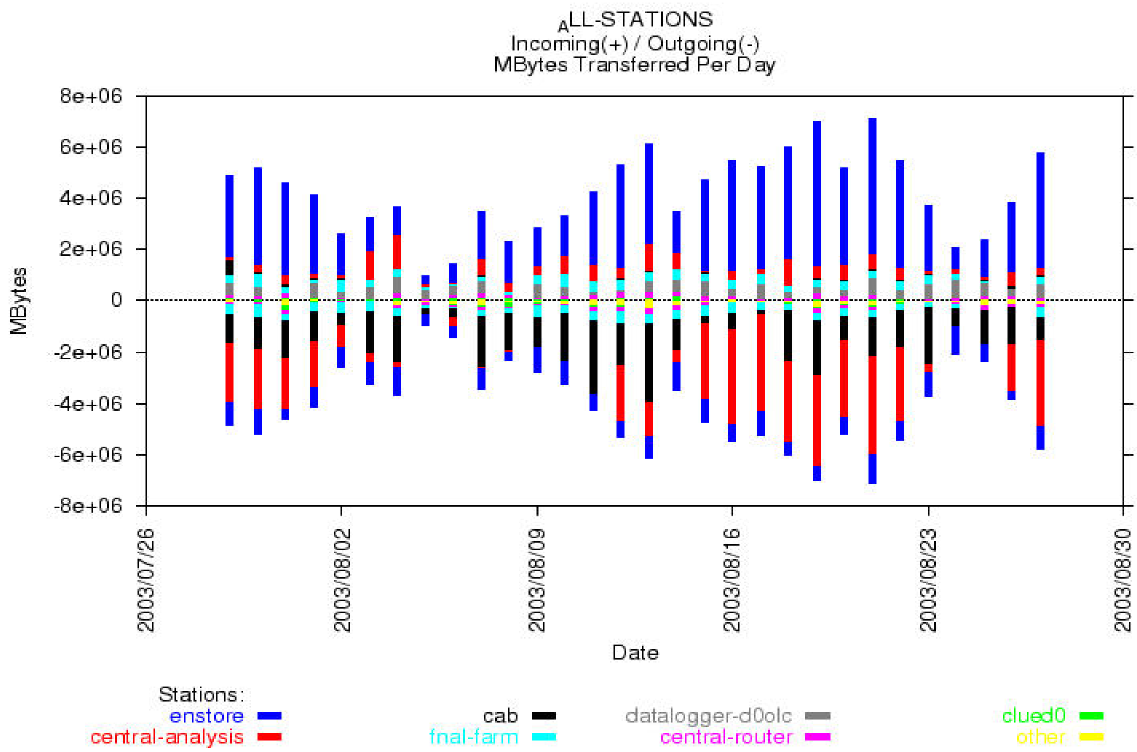

Figure 3.1 Data transfers in and out of all stations for a 30 day period. The colors correspond to different stations. The peak is 8 TB in/out in a 24 hour period.

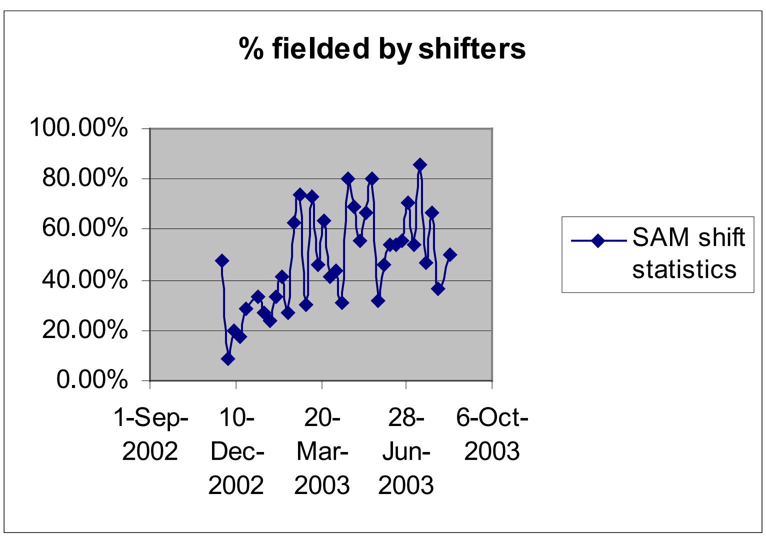

Figure 3.2 SAM shifter problem resolution as a function of time. This table tracks the resolution of problems that are sent in via e-mail. Shifters also use the monitoring tools to identify problems prior to user e-mail.

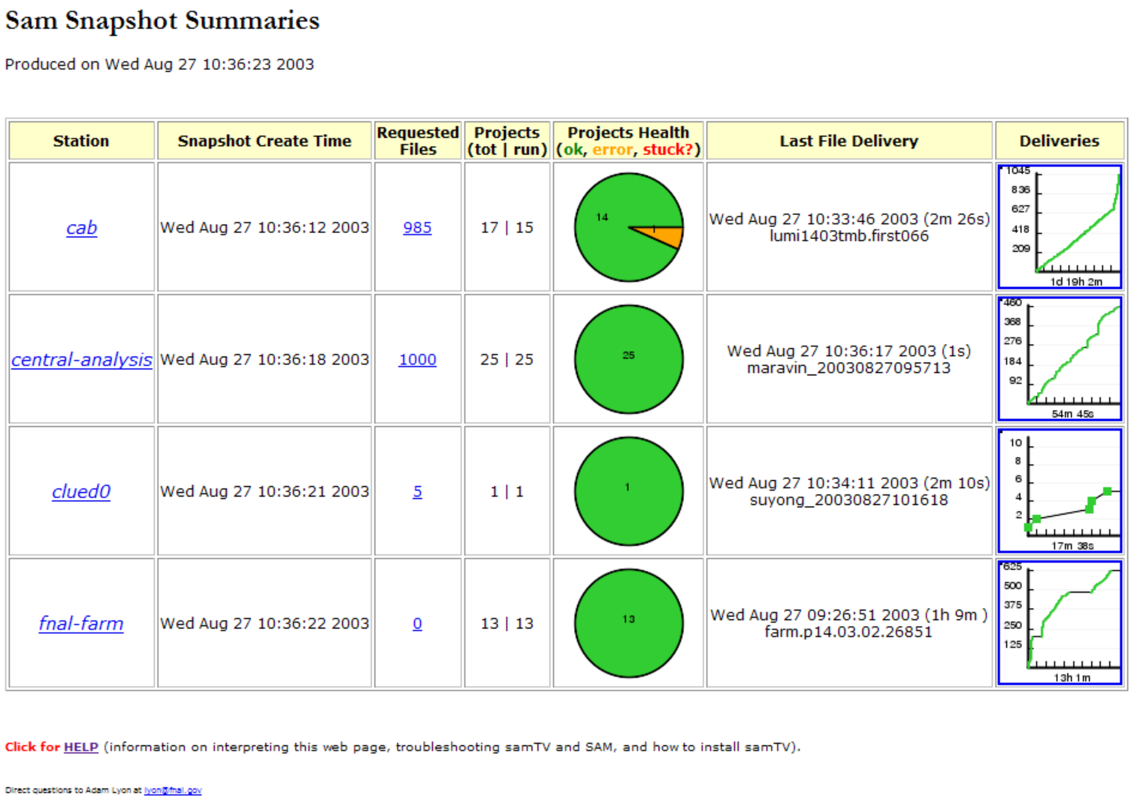

Sam Snapshot Summaries

Produced on Wed Aug 27 10:36:23 2003

Station	Snapshot Create Time	Requested Files	Projects (tot | run)	Projects Health (ok, error, stuck?)	Last File Delivery	Deliveries
cab	Wed Aug 27 10:36:12 2003	985	17 | 15	14 1	Wed Aug 27 10:33:46 2003 (2m 26s) lumi1403tmb.first066	1d 19h 2m
central-analysis	Wed Aug 27 10:36:18 2003	1000	25 | 25	25	Wed Aug 27 10:36:17 2003 (1s) maravin_20030827095713	54m 45s
clued0	Wed Aug 27 10:36:21 2003	5	1 | 1	1	Wed Aug 27 10:34:11 2003 (2m 10s) suyong_20030827101618	17m 38s
fnal-farm	Wed Aug 27 10:36:22 2003	0	13 | 13	13	Wed Aug 27 09:26:51 2003 (1h 9m) farm.p14.03.02.26851	13h 1m

Click for HELP (information on interpreting this web page, troubleshooting samTV and SAM, and how to install samTV).

Direct questions to Adam Lyon at lyon@fnal.gov

Figure 3.3 SAM-TV showing the health of four major FNAL stations. Shown are the number of requested files, the number of projects running, the pie charts displaying delivery states for the projects, and the last file delivered with the delivery rate. This is the top level of monitoring and more detailed information is available by clicking on some of the topics.

As CAB is the primary analysis platform, most event consumption occurs there. On DØmino much of the activity is driven by selection of limited numbers of interesting events from raw data tapes. This "picking" leads to large amounts of data consumption (comparable to CAB) on relatively fewer events. CLuEDØ, as a desktop cluster, is typically used to test and debug code on small samples. Most of the analysis at GridKa is based on skimmed data samples selected by the physics groups. As the diversity of usage suggests, these four stations provide a comprehensive look at the analysis systems.

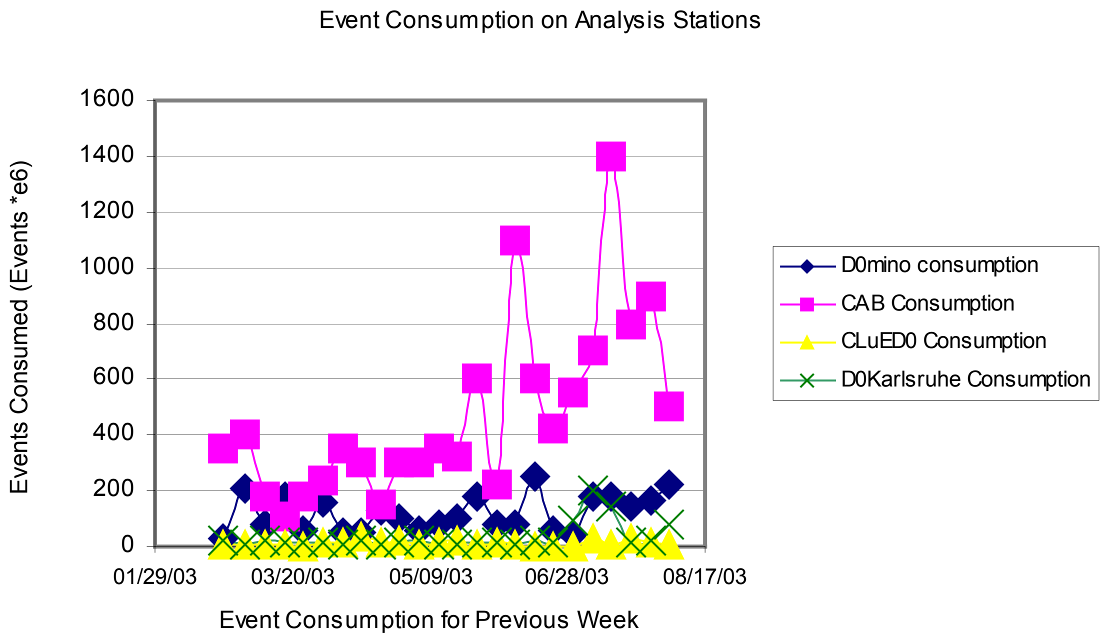

Figure 3.4 shows the number of consumed events on four of the analysis stations. Central analysis is shown in blue, CAB in pink, CLuEDØ in yellow and DØKarlsruhe in green. As CAB is the primary analysis platform, most event consumption occurs there.

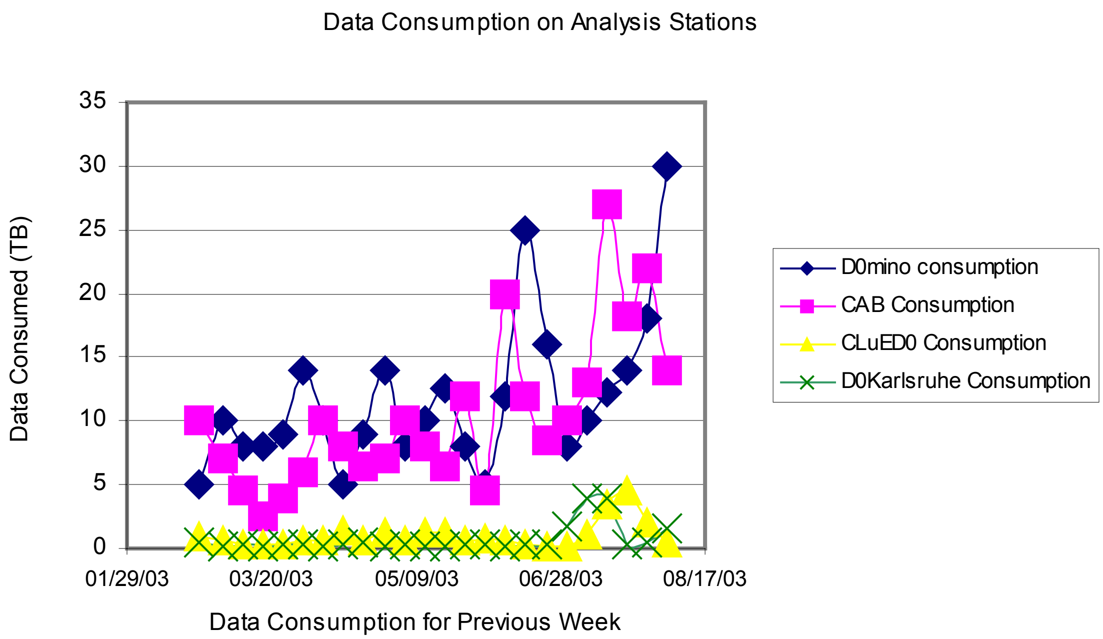

Figure 3.5 shows the amount of consumed data in terabytes on four of the analysis stations. Central analysis is in shown in blue, CAB in pink, CLuEDØ in yellow and DØKarlsruhe in green. Comparing this plot with Figure 3.4 shows these platforms fill different niches.

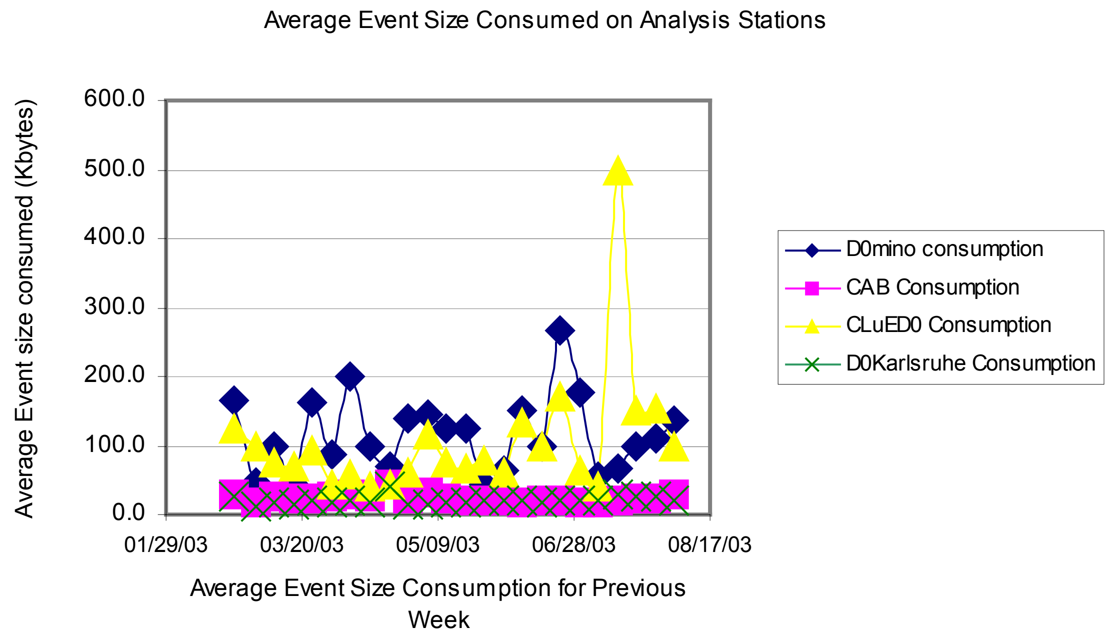

Figure 3.6 shows the average data size in kbytes on four of the analysis stations. Central analysis is shown in blue, CAB in pink, CLuEDØ in yellow, and DØKarlsruhe in green.

Figure 3.6 shows that CAB and GridKa event sizes are quite small, which indicates that thumbnail analysis is the primary type of analysis on CAB and at GridKa. The slight event size increase on CAB in August corresponds to the processing of raw data and DST data for reconstruction and for data quality studies. The event size on DØmino varies more widely than on CAB or GridKa. There are "legacy" analyses on DØmino by physicists who are completing work with old reconstruction versions; such analyses typically use the nearly obsolete 80 Kb/event root-tuple format. All of the pick-event activity currently occurs on DØmino through accessing of raw or DST data. CLuEDØ demonstrates the most widely varying data size. The one week spike of 500 Kb/events seen on CLuEDØ resulted from the study of special commissioning data. File transfer information is available for the stations as well, but is not shown here.

3.1 Evolving the SAM system into the Grid

SAM, as a fully distributed system, allows readily for integration with other Grid components. There already exists a prototype component of SAM for Grid job submission and monitoring (JIM), which is being tested and which is targeted to be in use for production simulation job submission by the end of 2003. JIM uses standard grid middleware (Globus and Condor tools). With JIM, plus the distributed components of the SAM core software, we can treat DØ-owned systems as a Grid. We also plan to go beyond that situation, as the SAM project continues to improve and augment its components, to use the SAMGrid to access general Grid systems in more transparent ways and to make the DØ Grid more robust and broadly useful to the collaboration. The areas of work to make these advances possible involve: (1) incorporating virtual organization technologies from Grid projects, (2) continuing to keep up with new developments in the middleware we use, (3) using Grid strategies (and, as they become available in production-quality releases, Grid tools) to move executables, run-time environments, and job output files around, and (4) evolving SAM's ability to utilize different input and output caching systems. In all of these areas, the SAM project as implemented for the Run II experiments has much to contribute in terms of defining realistic requirements and timelines for these Grid developments. DØ effort is contributing to the development of an integrated production system for our simulation jobs, which includes work listed under (3) above. The SAM project itself is making progress in (2) and (3), and will work on (1) and (4) during the next calendar year. We will also work on augmenting the successful SAM operations model to accommodate the greater complexity of Grid operations, and on adding the tools necessary to make the commissioning of general Grid operations possible with reasonable commitments of effort. In addition to the effort available for augmenting SAM, other DØ institutions have access to other resources, such as from LCG, and our direction is to make use of such resources. The DØ timescale for general Grid operations is approximately 2 years.

CHAPTER 4 – Disk and Robotic Storage

In this chapter, we describe our needs for disk, cache, tape drives and robotic storage.

The DØ data management system relies heavily on Hierarchical Storage Management (HSM) systems for archival storage. The principal HSM used by SAM is Enstore, developed at Fermilab by the Integrated Systems Department of the Fermilab Computing Division (now within CCF), and largely influenced by DØ requirements. Enstore is deployed at Fermilab and Lancaster University. SAM has also recently been interfaced to an HPPS storage system at the IN2P3 center. At FNAL, DØ has access to an ADIC AML/2 robot with LTO drives and an STK Powderhorn Silo with 9940A and 9940B drives for data storage. We have approximately 500 TB stored on tape.

The CCF department worked collaboratively with DESY to provide a disk cache and buffering system (dCache) that acts as a front-end buffer to the tape robot. This will provide direct interfaces to the cache through standard protocols like ftp and GridFTP, allowing any SAM station worldwide to access data directly from a dCache server without going through a specially configured Fermilab SAM station. Additionally, data that are being stored on tape will be available on disk for a short while for reading, allowing the reconstruction farm and the analysis jobs access without tape mounts. dCache is in production for CDF, and being interfaced to SAM. As seen in Figure 3.1, the current throughput for SAM can be as much as 8 TB in a day, and that seems a reasonable estimate for dCache disk cache. In FY2003, we purchased supporting hardware for dCache including a head node, logging node, and backup node, as well as two 2.5 TB file servers for tests.

To make an estimate of tape and disk storage needs, we identified possible data tiers, and estimated event size for those data tiers. The total amount of storage will depend on the number of events collected. We also assign a factor representing the number of times an event is stored for each data tier. For example, each raw event is stored once. We allow for some reprocessing storage in the estimate, and assume a large amount of derived data will be archived. The assumptions are shown in Table 4.1.

We intend to decommission DØmino by the end of 2004. Accomplishing this will mean replacing DØmino's large cache, project disk and disk serving capability with terabyte file servers. We intend to keep an SGI machine to serve the user home areas to CLuEDØ, the desktop cluster.

4.2 Storage Assumptions

		size	tape factor	disk factor
rates	average event rate	16 Hz		
	raw data rate	5 MB/s		
	Geant MC rate	3.2 Hz		
sizes	raw event	0.25 MB	1	0.01
	raw/RECO	0.5 MB	0.2	0.01
	data DST	0.2 MB	1.5	0.3
	data TMB	0.025 MB	3	1
	data root/derived	0.04 MB	9	1.5
	MC DØGstar	0.7 MB	0.1	0
	MC DØSim	0.3 MB	0	0
	MC DST	0.3 MB	1	0
	MC TMB	0.02 MB	3	0.2
	PMCS MC	0.02 MB	2	0.5
	MC rootuple	0.02 MB	0	0

Table 4.1 Event size and stored data for tape and central analysis disk cache is shown. The columns labeled "tape factor" and "disk factor" show the fraction of events on tape and disk for each tier relative to raw data. The above tape and disk factors should be taken as representative—different assumptions apply to the FNAL virtual center, the FNAL realized center, and different remote centers. These tables are directly from the planning spreadsheet, which has some artificial distinctions between formats.

The tiers are raw/RECO (RAW), data summary (DST), and thumbnail (TMB). The raw/RECO (RAW) data tier includes the raw data and the reconstructed output. These samples are useful for trigger and reconstruction studies and those analyses which need more information than the DST provides. The data summary tier is expected to have sufficient information to allow for reprocessing. We assume that more DST than RAW files will be stored to tape to allow for reprocessing. The thumbnail is a physics summary format, and is presumed to be the starting point for most physics analyses. We assume that DST-level reprocessing will produce additional TMB copies that must be concurrently stored. We anticipate that most derived data sets will be subsets of the thumbnail. Based on Run I history and reinforced by our Run II experience we allow for a large amount of these sets to be stored on tape.

The amount of Monte Carlo tiers which must be stored are subject to trade-offs among tape costs, the need to re-reconstruct and to re-run the trigger simulation, and to simulate different instantaneous luminosities. We prefer to assume a generous amount of MC DST storage. We assume that there is one primary TMB sample on disk on the central analysis system and the derived data sets are kept on physics group project disk servers. With these assumptions, the FNAL storage needs are shown in Table 4.2, with 750 TB of tape storage and 85 TB disk cache and project disk available per data-collection year. The tape cost under these assumptions is approximately $280,000 per year given a tape cost of $0.35/GB, and corresponds to 4000 tapes. DØ could recycle roughly 2000 tapes in FY2004. DØ plans to migrate to a new tape or drive technology on the time scale of 2006. Moving to 400 GB/tape drives in the near term may not be cost-effective for DØ. We have a significant capital investment in 9940B drives and intend to run them in production through 2006, with 2006 as a transition year, and 2007 targeted for a full migration. Our recent experience is that it takes roughly nine months from the delivery of drives for beta testing until they can be used for production. Our more distant experience is that some commodity drives fail the testing and cannot be used for production. On that basis, migration would next be cost effective when 800 GB/tape media is available. We budget in 2006 to begin to acquire such drives.

data samples (events)

	1 day	1 year
event rate	1.38E+06	5.05E+08

TAPE data accumulation (TB)

raw event	0.35	126.14
raw/reprocessing	0.14	50.46
data DST	0.41	151.37
data TMB	0.10	37.84
data rootuple	0.50	181.65
MC DØGstar	0.10	35.32
MC DØSim	0.00	0.00
MC DST	0.41	151.37
MC TMB	0.08	30.27
PMCS MC	0.06	20.18
MC rootuple	0.00	0.00
total storage (TB)	**2**	**785**
total storage (PB)	**0.002**	**0.78**
total storage (GB)	**2,150**	**784,616**
total storage MC (GB)	**1**	**237**

	1 day	1 year
event rate	2.16E+06	7.88E+08

DISK data accumulation (TB)

raw event	0.00	1.26
raw/reprocessing	0.01	2.52
data DST	0.08	30.27
data TMB	0.03	12.61
data rootuple	0.08	30.27
MC DØGstar	0.00	0.00
MC DØSim	0.00	0.00
MC DST	0.00	0.00
MC TMB	0.01	2.02
PMCS MC	0.01	5.05
MC rootuple	0.00	0.00
total storage (TB)	**0**	**84**
total storage (PB)	**0.000**	**0.08**
total storage (GB)	**230**	**84,012**
total storage MC (GB)	**0**	**7**

Table 4.2 shows the total data storage required for assumptions listed for one year

The assumptions for the amount of disk storage were based on recent work by a committee charged with looking data set definitions with respect to the physics needs of the experiment. The total disk storage corresponds to about 80 TB/year at FNAL. 30% was added for contingency, and 10% (about 8 TB) set aside for dCache. Eventually, within the virtual center, we might like to have all DSTs on disk—but this would require working Grid applications, and further, there is still some debate about the need for having all DSTs on disk. We assume that the current cost of disk servers is $20,000 for 5 TB of disk and assume a doubling every 18 months.

File Server Cost Estimate

cost/fileserver	10,000
Network cost/16 FS	10,000
Contingency	40%

Year	Capacity(TB)
2003	2.5
2004	3.5
2005	5.5
2006	8.7

Data Volume	FY04		FY05		FY06	
84.01	No. FS	Cost	No. FS	Cost	No. FS	Cost
	33	360,000	21	230,000	13	140,000

Table 4.3 shows the cost for analysis and dCache file servers. The contingency corresponds to 30% on the disk estimate and 10% to account for the dCache servers.

The current drive usage is shown in Figures 4.1-4.3. As we are decommissioning the 9940A drives, the usage is tailing off, although it will pick up again as we begin the reprocessing. The usage on the 9940B drives is extensive, with all 16 drives often in use. We are not yet using these drives at full rate, and intend that the introduction of dCache will mitigate the need for more drives both by enabling DØ to operate the drives at rate, and by reducing the number of drives in use by providing a tapeless data path. The LTO drive usage is also shown, and we are in the process of putting LTO2 drives in production.

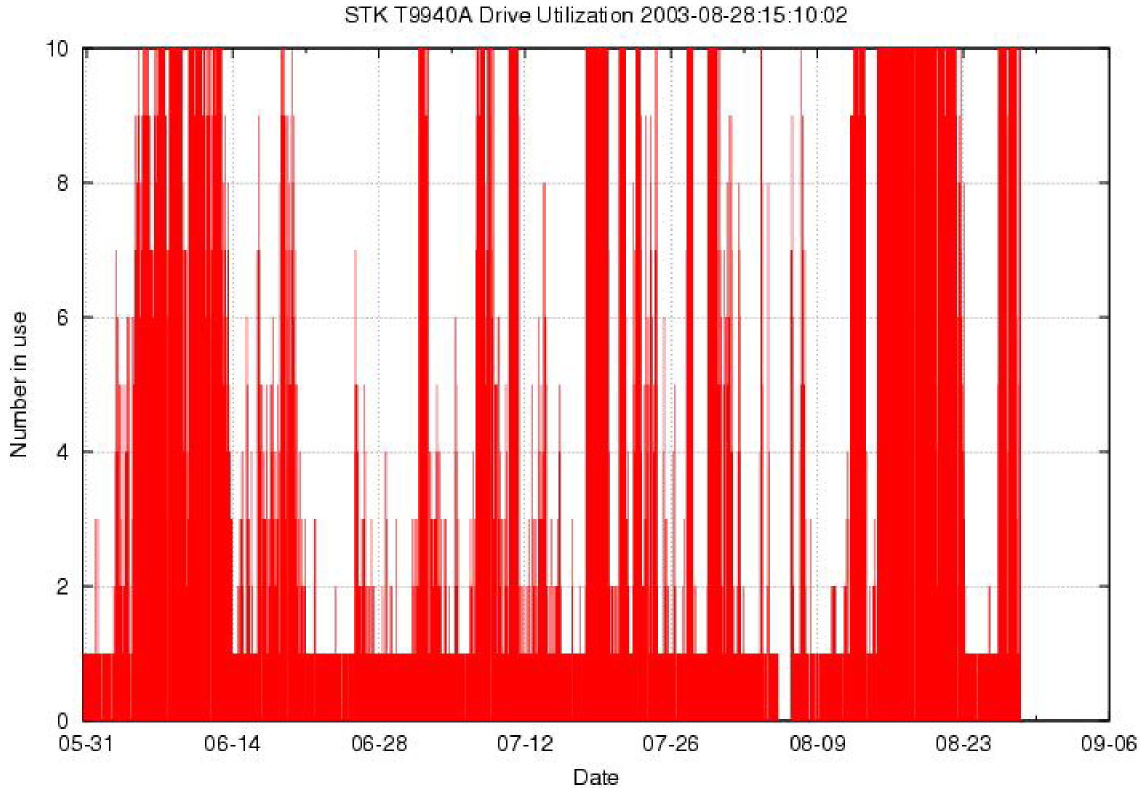

Figure 4.1 shows the load on the nine existing 9940A drives. Usage is tailing off due to the fact that all recent data are on the 9940B or LTO drives.

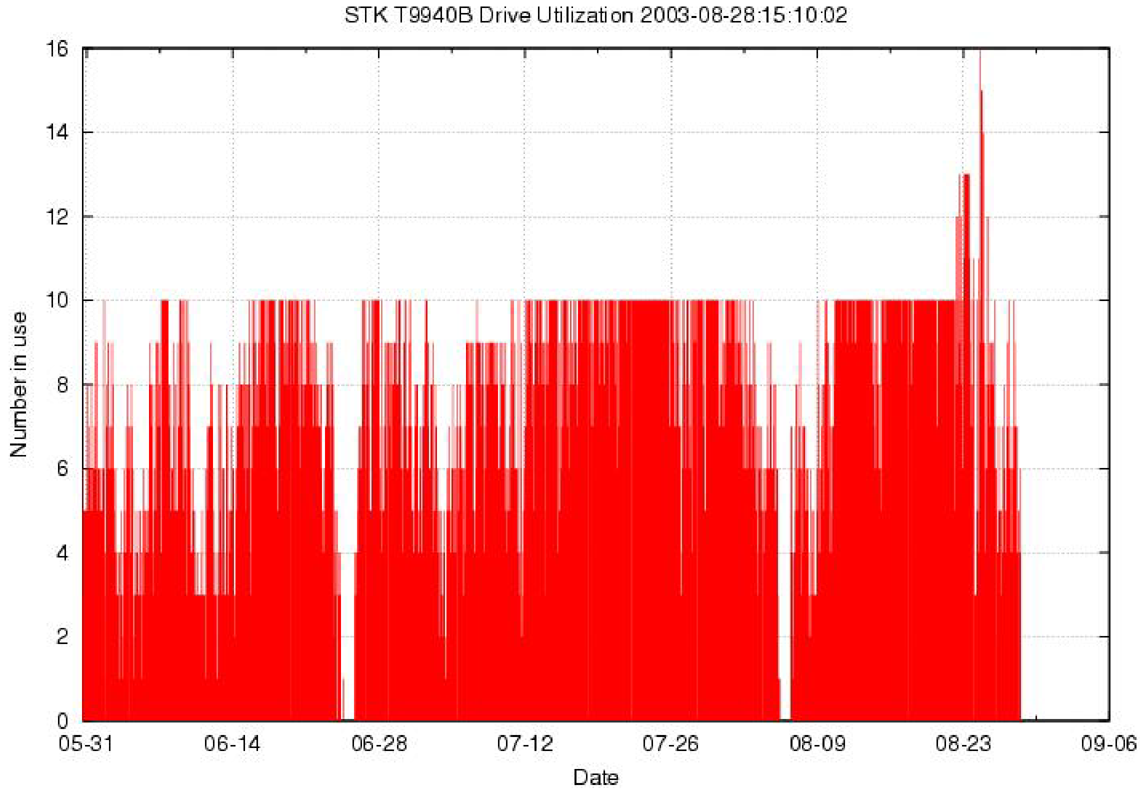

Figure 4.2 shows the load on the 16 existing 9940B drives. Production for 9940B drives began in Febuary 2003 with reconstructed output, and in May 2003 for raw data. 6 new drives were recently put in service and 4 more will enter service soon. More recent plots show all 16 drives are often used, due to significant tape access from the selection of individual event samples.

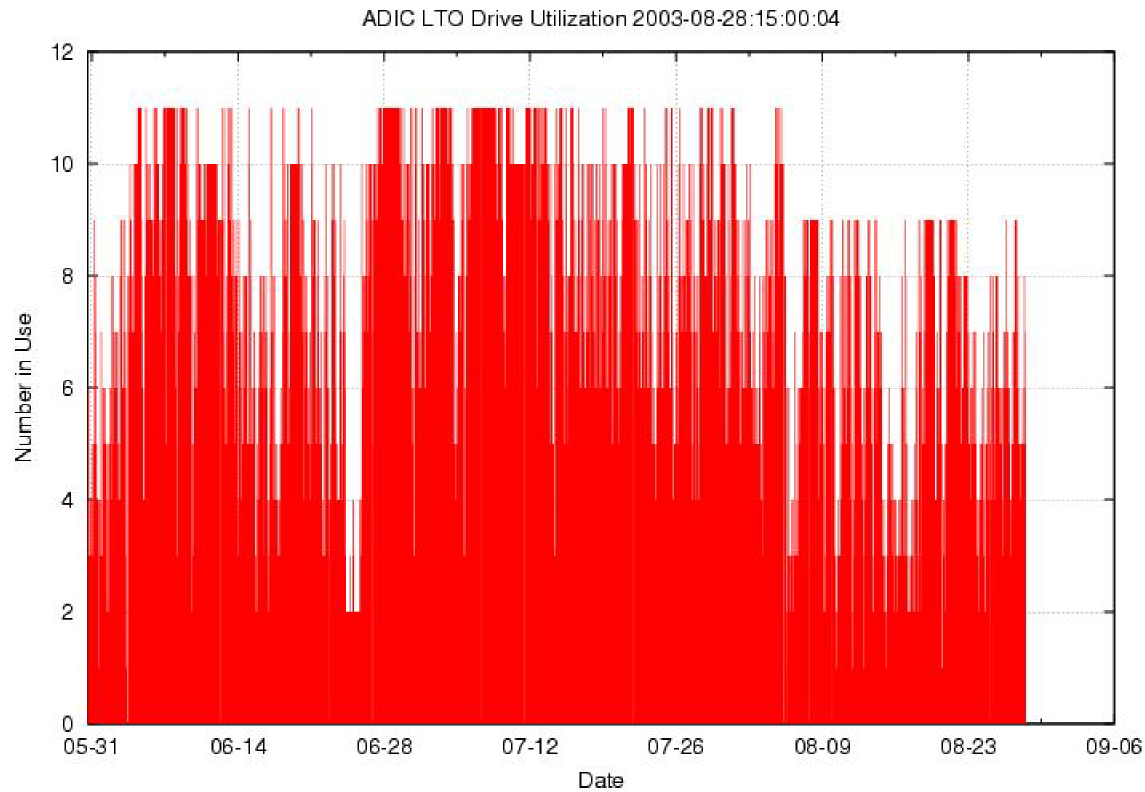

Figure 4.3 shows the load on the 9 LTO drives. Monte Carlo data and TMB files are stored in the ADIC robot. The other 2 drives are LTO2 drives which are being commissioned with a 30 TB test.

Each 9940B drive costs \$30,000, and each LTO2 drive costs \$8000. In addition every drive requires a \$4000 Enstore mover node. In 2004, we plan to purchase 5 9940B drives and 5 LTO2 drives for a total of \$230,000 in drive purchases. After that point, we would probably chose to add more dCache nodes rather than more 9940B drives. A \$100,000 expansion of LTO2 drives is planned for FY2005. In FY2006, we allocate \$500,000 to begin to replace the 9940B drive plant with higher density drives and media.

CHAPTER 5 – Production Computing Systems

This section describes the large-scale systems used for production processing, simulation, and data analysis. These systems are located at the Feynman Computing Center (FCC), the DØ experimental area, and worldwide.

5.1 Local Farms

The current DØ production farm system employs 122 dual-processor Intel worker nodes and an SGI O2000 used as an I/O server node. Recent performance of the farm is shown in Figure 5.1. The total capacity of the system is approximately 800 GHz. It is expected that this will be expanded in the next few months with 96 dual processor 2.6GHz nodes. This will bring the total capacity to 1100 GHz equivalent. When this happens, it is expected that some of the slower current nodes will be retired, and the 1 GHz nodes will be dedicated to input file staging. All worker nodes are connected via 100Mb/s interfaces to the same Cisco 6509 switch as the I/O nodes.

For future planning, we assume that the farm systems are expected to remain operational regardless of accelerator shutdowns. The 2003 farm purchase found the best price/performance point was to purchase 2.6 GHz machines. Recent historical trends have seen new generation systems introduced at a roughly constant $2200 per dual processor unit. We have assumed that machines will need replacement once they reach three years of age or more. We also assume that the I/O nodes will have to be scaled up to handle the increased load.

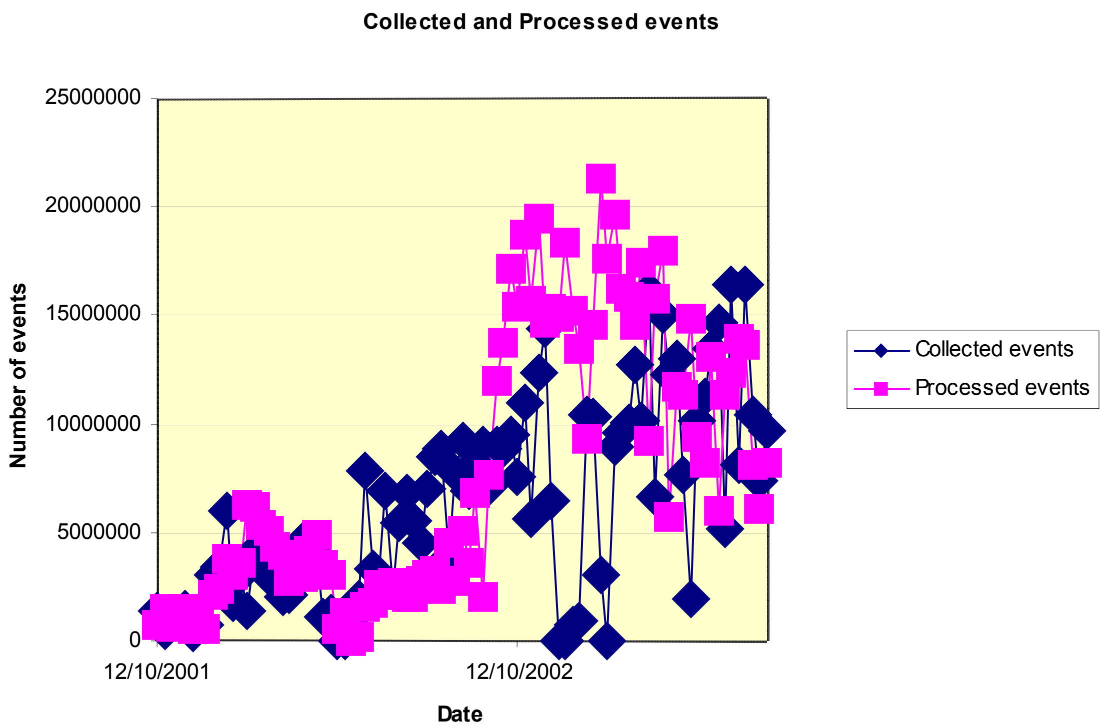

Figure 5.1 Central farm performance per week (in pink) and the number of collected events per week (in blue). The January shutdown is visible for the collected data, and corresponded to a very productive period for farm production. The turn-on of the 2002 purchase can be seen in Nov '02. In March '03, increased instantaneous luminosity led to decreased throughput for the farms (see Figure 2.1), and in July '03, operational problems with p14 hampered throughput.

All accounting numbers for the farm processing information are stored in SAM, and the parentage of a file is traceable back to the raw data file from which is it derived. That enables us, for example, to check the ratio of the data delivered to user analysis compared to the raw data that has been collected. For the p13 version of the reconstruction, this was found to be 98.2% for the W/Z group skim. The percentage of data processed directly with p13 is summarized in Table 5.1. Data prior to September 2002 were not fully reprocessed due to detector problems. Some data in November 2002 – January 2003 period were processed multiple times, also due to detector or reconstruction problems. The losses in this table include those due to data handling and reconstruction failures (the latter is by far the largest source of failure). As can be seen, the production chain is working well.

Date	Raw		% proc >=13.05	
	Files	Events	Files	Events
May-02	7796	1.41E+07	0.76	0.74
Jun-02	1713	2.93E+06	0.30	0.32
Jul-02	11041	2.17E+07	0.00	0.00
Aug-02	11247	2.12E+07	0.54	0.56
Sep-02	14456	2.89E+07	0.98	0.98
Oct-02	15741	3.51E+07	0.97	0.97
Nov-02	16458	3.56E+07	1.09	1.10
Dec-02	17223	3.79E+07	1.37	1.38
Jan-03	6783	1.82E+07	1.60	1.62
Feb-03	7440	1.94E+07	0.99	0.99
Mar-03	11783	2.96E+07	0.96	0.96
Apr-03				
	121681	2.65E+08	0.93	0.96

Table 5.1 shows the number of files and events collected and the percentage processed with p13.

5.2 Remote Facilities

Over the next few years DØ will rely on remote computing facilities for a significant fraction of its processing power. There are three major tasks we see being carried out by these sites:

- Monte Carlo (MC) Production
- Secondary reprocessing of the data.
- CPU intensive user analysis jobs.

Each of these modes of operation imposes its own requirements. The simplest is MC production since it is essentially self-contained and does not require access to calibration or other databases. Reconstruction processing of data requires database access and also needs more careful bookkeeping for specific binaries. User analysis jobs have diverse input data, output data, and application requirements.

The collaboration has made a survey of available current and projected large scale computing resources. The results are summarized in Table 5.2. Additional resources can often be obtained on an as-needed basis or as opportunities arise. The bare minimum required capability of these production facilities is that they be sufficient to meet all of the needs for MC production. However, the DØ physics program will be significantly enhanced if we have sufficient resources to allow reprocessing of data and large-scale remote analysis, and this is our policy and intention. It is understood that these remote production facilities are in many cases shared with other experiments and are not able to upgrade operating systems purely to meet DØ software requirements.

Our approach in planning and costing for computing has been to estimate the total requirements as if they were all placed at Fermilab. Remote facilities that are made available to DØ then receive credit to DØ's Operations Common Fund in reflection of the savings that they generate. The total costs quoted in this document therefore do not reflect either the expected actual cost to Fermilab or the expected actual cost at the remote installation, which may be subject to many local factors. They reflect the cost of a "virtual center" at Fermilab that will never actually exist, but if it did, would be able to carry out all of our computing needs locally.

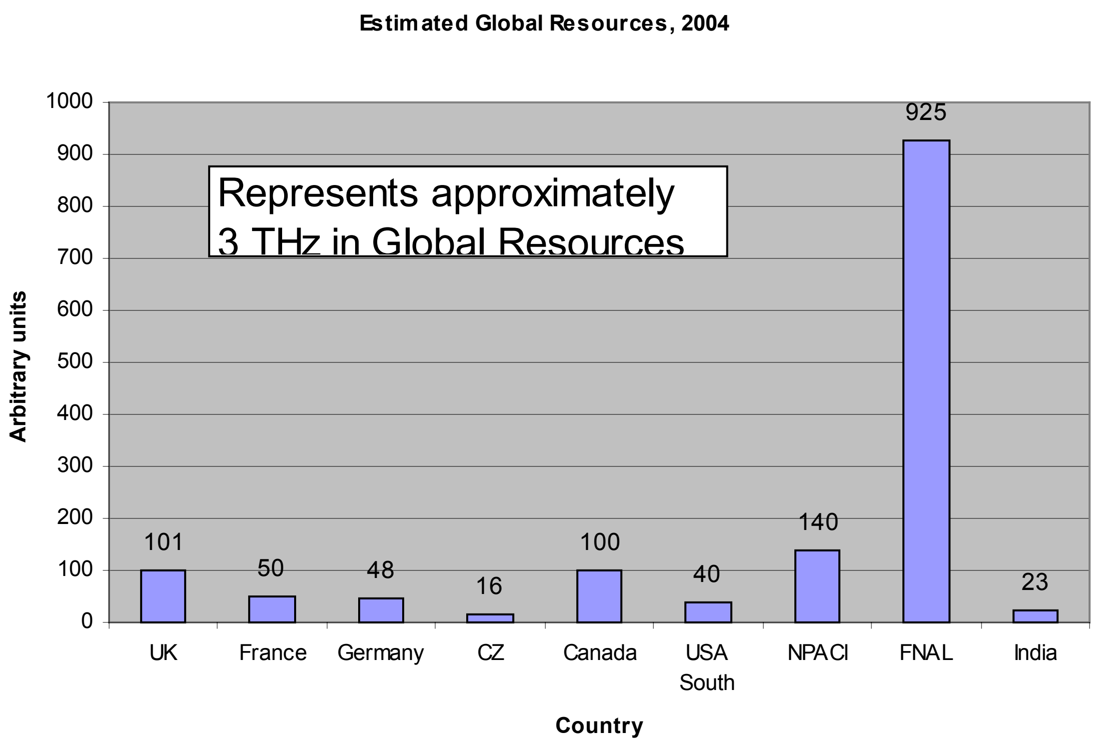

*Figure 5.2 Current understanding of the global computing resources available to DØ, in 2004. Our understanding evolves, not all resources are represented and some resources may be over counted. The units are SpecInt2000*1000, and correspond to 3 THz. For example, the NIKEHF resources are not shown. Additional resources can often be obtained on an as needed basis or based on opportunity. However, this plot does demonstrate that there is approximately 1:1 ratio of FNAL resources to non-FNAL resources.*

All Monte Carlo production for DØ has taken place remotely. Figure 5.2 shows the number of event in millions that have been produced. As can be seen the slope has remained basically constant, with a slight decrease this year. This is a success story; however more resources will be needed as we enter the publication phase. DØ treats Monte Carlo as a production activity with all of the requests coordinated centrally, and the farms use common workflow software called mc_runjob, which, among other tasks, is responsible for producing the SAM metadata. Currently, all MC data are stored at FNAL in the ADIC robot. One of our goals for the next year is to further automate the submission of Monte Carlo jobs using Job and Information Monitoring (JIM) and possibly storing some of the MC files at remote sites.

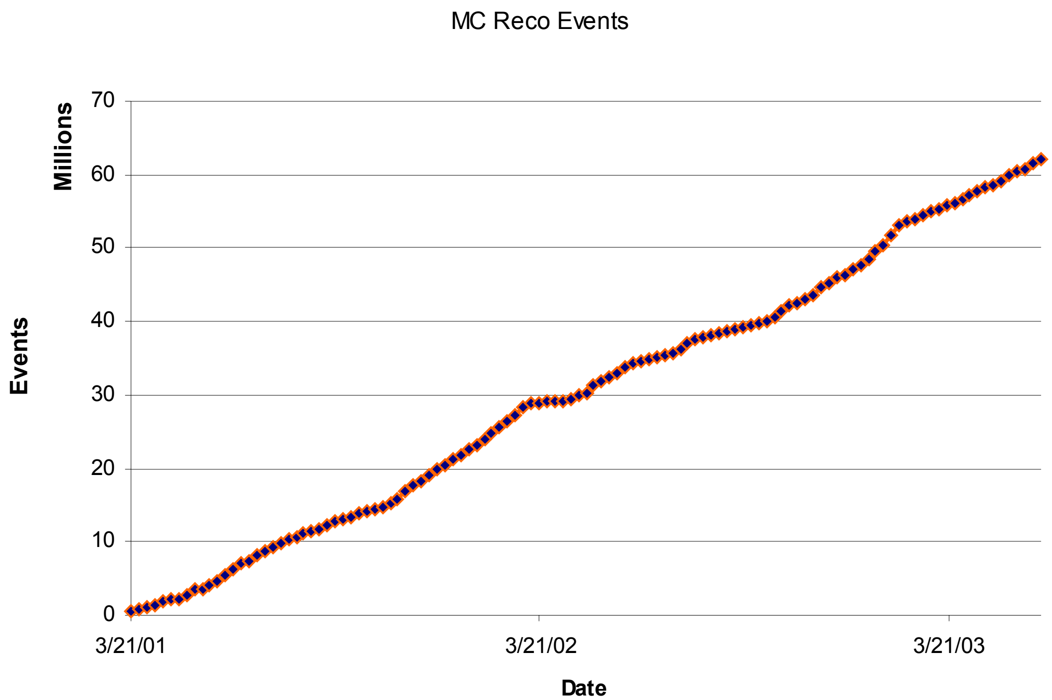

Figure 5.2 Monte Carlo production since March 2001.

Planning and testing has started for a Fall 2003 reprocessing of a partial data set using global resources, which are typically not dedicated to DØ. Runjob, an extension of MC_runjob, will be the default script for the reprocessing. This project is intended to provide improved p14 tracking to analyses targeted for publication this winter. The collaboration has explored a number of the scenarios for reprocessing using the available resources shown in Figure 5.2. We determined that 150 million events should be selected for reprocessing in a 90 day period.

5.2 Cost Estimates for Production Systems

Three RECO time/event possibilities are shown. At this time, we anticipate 50 sec/event for the 2004 instantaneous luminosity guidance. We show primary reprocessing costs assuming 50 GHz-sec/event in 2004, 60 GHz-sec/event in 2005 and 80 GHz-sec/event in 2006. We assume that the local FNAL farm efficiency is 70% , and that the online rate averages 16 Hz, with a peak rate of 50 Hz to tape, and 30% combined accelerator duty factor and DØ data collecting efficiency. The average of collection rate of 16 Hz is in agreement with the measured collection rate over the past six months.We assume 3 GHz processors in 2004, with 6 GHz processors available in 2006, with a fixed price per machine (dual processors) of $2200, and an I/O cost of $25,000 assigned to each 100 nodes. We take into account the existing plant and plan to retire nodes after 3 years when they leave warranty. For reprocessing we assume that the reprocessing centers are 50% efficient (as it is much harder to achieve high efficiency on shared resources), and that the reprocessing has a 90 day duration to reprocess 50% of the data collected the year before. We do not assume any legacy machines in this case, and thus after 3 years, the resources would be assumed to be able to process somewhat less than 50% of the full data set collected, as we expect that the reconstruction gets slower due to increased instantaneous luminosity. The 90 day duration for reprocessing is driven by the need to have a reasonable development cycle for improvements to the reconstruction, to avoid long reprocessing with obsolete reconstruction versions and to accommodate student schedules.

For simulation, our goal is to generate about three Monte Carlo events for every four collider events collected. Using the same assumptions as in the farm production profile, Table 5.4 shows the cost and number of nodes which the regional centers would have to purchase to meet this need assuming a mix of plate level (detailed) and fast simulation. In the detailed simulation, each event is overlaid at the digitization stage by with zero-bias events (random sample of the detector) to simulate noise and additional soft interactions. An average of 170 seconds corresponds to roughly one-quarter of the events using full simulation and the other three quarters using fast simulation.

Average Rate:	16
Farm Efficiency:	70%
Misc. Processing:	20%
Reprocessing:	0%
Cost/node:	2,200
I/O Cost/100 nodes	25,000

CPU	SpecI2000
2.6Hz	720
3GHz	960
4GHz	1280
6GHz	1920
10GHz	3200
15GHz	4800

2001	#of nodes	GHZ
2002	260	699
2003	96	258
2004	174	701
2005	181	974
2006	127	1,022

Primary Reconstruction Cost Estimate

year	2004	2005	2006
Reco time	50	60	80
Required CPU	1371	1646	2194
Existing system	670	672	1173
Nodes to purchase	174	181	127
Cost	$407,506.67	$423,464.00	$303,572.80
#Nodes at FCC	530	451	482

Reconstruction Cost Estimate

year	2004	2005	2006
reco time	50	60	75
duration	90	90	90
fraction	50%	50%	50%
Rate	32.44	32.44	32.44
Farm eff.	50%	50%	50%
#nodes	804	724	603
CPU required (GHz)	3244	3893	4867
	$ 1,969,574	$ 1,767,617	$ 1,477,181

Monte Carlo Cost Estimate

Year	2004	2005	2006
MC time	170	170	170
duration	275	275	275
fraction	15%	5%	5%
Rate	3.19	1.06	1.06
Farm eff.	70%	70%	70%
#nodes			
CPU required (GHz)	774	258	258
	$ 446,940	$ 105,485	$ 70,323

Table 5.3 Resources needed for production, including primary processing, reprocessing and MC production. The reprocessing of MC events is included in the estimate.

Using Figure 5.2 as a working estimate, we have roughly 1.5 THz of computing available to DØ in 2004. Not all of those machines can be used for production due to network limitation, and not all of those cycles are available at all times or for MC production.

5.3 Central Analysis Computing

The legacy central analysis system, DØmino, is an SGI Origin 2000 system comprising 128 R12000 (300 MHz) processors, with an attached data cache of ~ 50 TB fibre channel disk, with RAID disk for system needs, and user home areas. DØmino currently provides a centralized, stable, and uniform work environment with interactive and batch services for on and off-site users. DØmino also provides data movement into a cluster of Linux compute nodes, called the Central Analysis Backend (CAB). The configuration of the system includes 16 1.1 GHz dual nodes and 160 2.0 GHz AMD nodes, with an equivalent power of 500 GHz.

Usage for the CAB system is shown in Figures 5.3-5.5, which show the CPU time and wall time history and their ratio. This provides one measure of the effectiveness of SAM deliveries on the system. There are a number of SAM queues, *SAM_HI* (since decommissioned), *SAM_LO*, and a non-SAM queue, *Medium*. In terms of CPU time, the SAM and non-SAM usage is comparable. The typical of use of the SAM queue on CAB is to run over the TMB data sets to produce skimmed samples, storing those skims into SAM, and then producing analysis specific user formats for desktop cluster or laptop analysis. The usage pattern of the *Medium* queue reflects an unanticipated need from last year's planning, namely user level Monte Carlo generation for testing. In addition to MC generation, the other typical use of *Medium* is generation of user level files when skims are not stored in SAM.

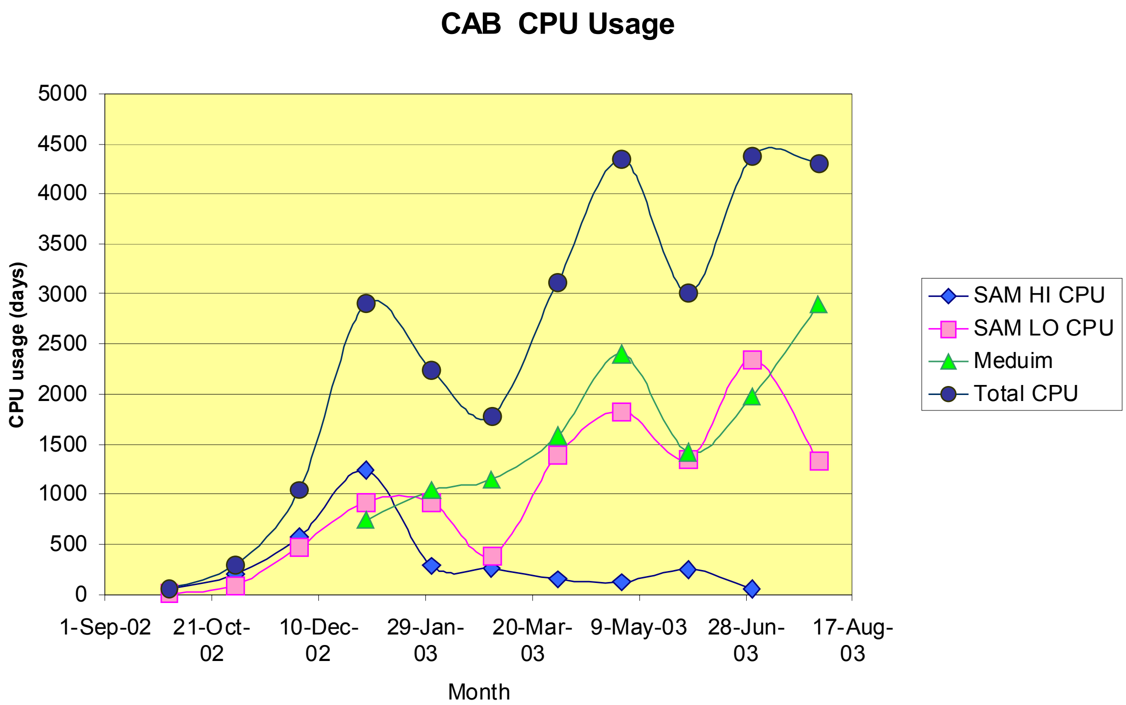

Figure 5.3 shows CPU time on the central analysis backend (CAB) for SAM_HI in light blue, SAM_L0 in pink, and Medium in Green. The sum is shown in dark blue. The queues are described in the text. August is a partial month.

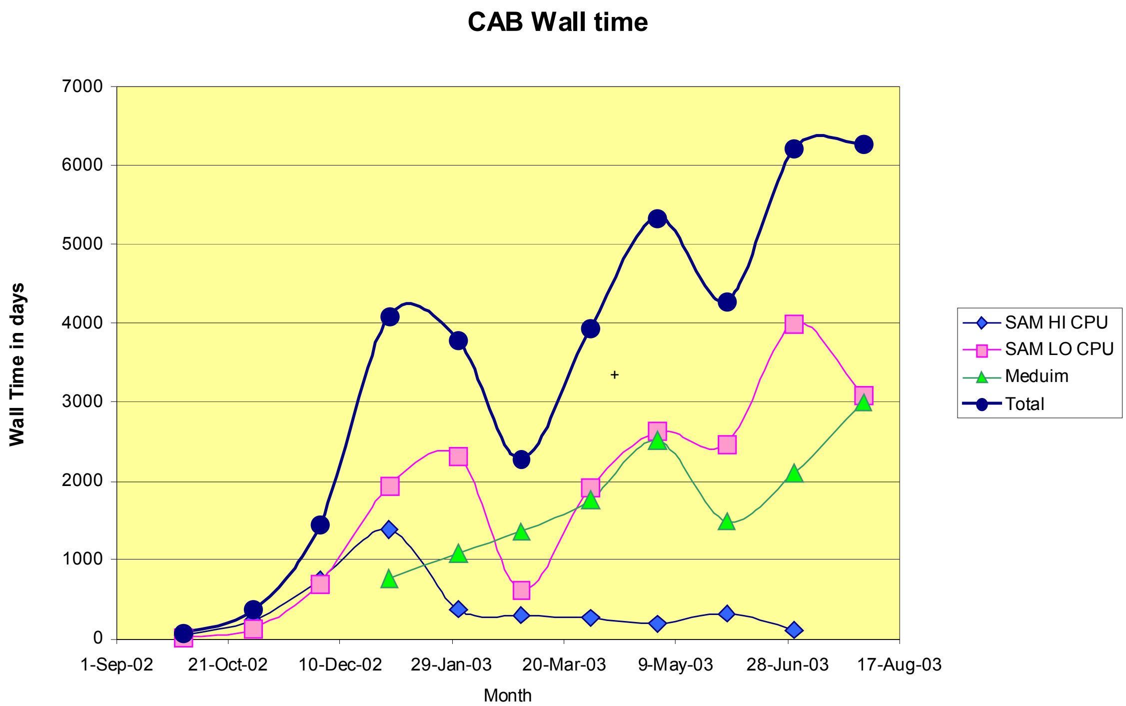

Figure 5.4 As above, but showing Wall time rather than CPU time.

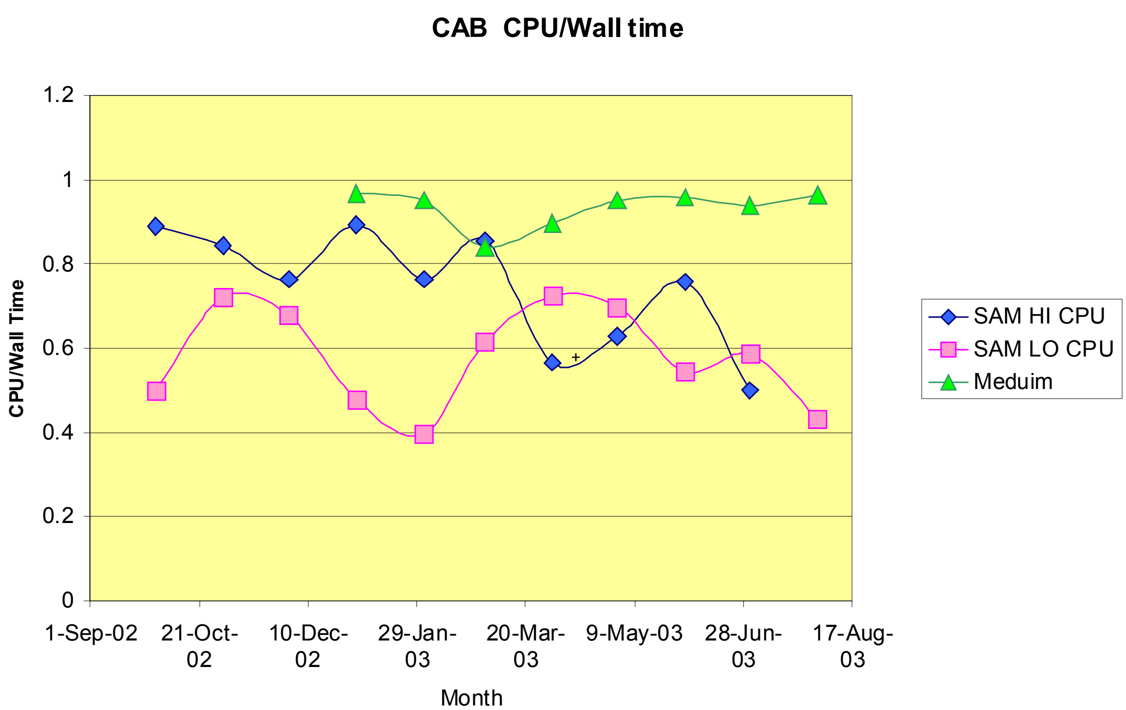

Figure 5.5 Ratio of CPU/Wall time on the central analysis backend (CAB). The queues are described in the text. The dips in SAM_LO correspond to heavy data analysis periods.

For the ratio of CPU to Wall time, the target for the SAM queues is about 70%. Two dips that represent performance significantly below that are visible in late January 2003; corresponding to the winter conference analysis period. In light of this experience, the performance of SAM on CAB was analyzed and improved, as can be seen for the April-June period. (The systems were lightly loaded in March.) The next dip, in August, was towards the end of the summer conference period. However, the dip was traced to increased tape access from CAB which was causing contention between disk and enstore access. Parameter tuning mitigated much of the problem; however, other issues have also been uncovered and are being addressed.

For planning purposes, we have estimated the CPU needed for "large-scale" analysis processing, such as the type of data skimming and processing jobs currently run on CAB to prepare small samples for user desktop analysis. It does not include the CPU needed for the desktop analysis itself. We have also assumed that the CPU required will be proportional to the amount of data. Based on current CAB usage and our understanding of worldwide analysis needs, we estimate that each 500M events (1 year's data) generates an analysis need of 7000 CAB cpu-days/month. We assign a contingency of 100% to accommodate peak loads. The total analysis need then corresponds to 700 GHz, including remote analysis installations, but not including desktop clusters, as noted above. The estimated cost to supply this CPU is shown in Table 5.5 below. The cost and evolution of processor speed is also the same as in the farm estimate, as is the assumption that machines that are three years old will have to be replaced.

Offline Efficiency:	100%
Contingency:	100%

Calculated CPU with efficency										Total	
Analysis	THz CPU	FY03, 2.6GHz Nodes		FY04, 4GHz Nodes		FY05, 5GHz Nodes		FY06, 6GHz Nodes		Target	
	Per Year	No. Nodes	Cost	No. Nodes	Cost	No. Nodes	Cost	No. Nodes	Cost	No. Nodes	Cost
Analysis CPU	0.70	210	420,000	157	314,000	118	236,000	78	156,000	563	970,000
Replacement	0.00	0	0	0	0	118	236,000	78	156,000	196	236,000
Total to Purchase:	0.70	210	470,000	157	339,000	236	522,000	156	337,000	759	1,331,000
#Nodes At FCC		370		527		603		549			

Table 5.5 Cost estimate for "large-scale" analysis CPU, not including desktop cluster analysis.

5.3 Remote Analysis

Remote analysis is important because it will enable us to fully exploit the huge amount of data that we are collecting. By bringing the data to every one of the physicists in the collaboration, we can not only enhance their individual productivity and experience, but we can gain computing resources, and improve the overall physics output of the experiment. The DØ collaboration and its management are fully committed to the development, enhancement and wide use of the capability to carry out physics analyses remote from Fermilab. This desire reflects the internationalization of the collaboration over the past five years (more than half of the collaboration is now based at non-US institutions), recent technological developments, and the realization that this direction will be critical for the success of our own work at the LHC. Our remote analysis strategy builds on the success and far-sighted design of SAM, which provides the data backbone for the system, and our use of significant distributed computing resources. It does not, of itself, require extra Grid software; as explained in section 3.1, our strategy is to use what exists, align ourselves with ongoing Grid developments and to incorporate additional Grid tools as they become available. DØ is a member of the Open Science Grid initiative.

During 2002, an Offsite Analysis Task Force was set up and under its aegis a prototype Regional Analysis Center was established at GRIDKA in Karlsruhe. This center was used as the source of data for a successful remote analysis, a search for new particle production by a German group that was presented at Moriond 2003. Similar efforts are underway in many other collaborating nations in Europe, and also in a cluster of our universities in the southwestern US. While we are still early in this process, it has become clear that there are big sociological and organizational issues to be addressed if physics analysis at remote institutions is to become a mainstream activity. There are human and communication challenges in getting analysis systems and software set up and running at remote sites. Once these are addressed, there are ongoing challenges to be faced by those actually trying to "do physics" remotely. These include the nature of meetings, how informal and formal information flows are handled, documentation, results-approval procedures, and so on. The same challenges will be faced in the LHC era, no matter how polished the Grid software may then be. We think that DØ's experience will provide a useful starting point for understanding the problems, and how to address them. So far, we have adopted and customized the CERN Document Server for use as a (mandatory) DØ meeting agenda server and presentation material archive. (It is, naturally, hosted remotely, in Nijmegen.) Physics meetings have been rescheduled to take place only in the mornings, to allow videoconference connections to Europe. We are investing in a new, substantially enhanced videoconference room at DØ to improve the quality of connections. Analysis documentation requirements are being clarified, and so on. It is clear that making remote analysis straightforward and natural will take several years, but we are enthusiastic about doing so.

CHAPTER 6 Budget Summary

In addition to the costs outlined above, there are infrastructure costs associated with running the experiment. The primary sources of these are for database machines, disk and servers, networking, miscellaneous machines such as those used to build the DØ code base, and web servers. We also often have to pay for other infrastructure costs such as purchasing raid arrays for Enstore/dcache. Table 6.1 shows our estimated infrastructure costs.

	2004	2005	2006
databases			
servers	$30,000	$30,000	$30,000
disk	$30,000	$30,000	$30,000
Networking	$120,000	$80,000	$100,000
Machines	$60,000	$60,000	$60,000
Totals	$240,000	$200,000	$220,000

Table 6.1 Cost estimate for infrastructure.

The database disk and server estimates are based on experience and scaling the current usage. We purchased a new SUN machine in 2003, and do not anticipate another major upgrade of the database machine itself until 2007. "Machines" are build machines, web servers and other small requests that come up for servers or disk that can be capitalized to one of the major systems. We will also need to add faster disk to DØ2ka or use an alternative solution in order to serve home areas, and we allocate $50,000 for this purpose.

The 2004 networking costs are estimated to include a 6513 chassis for the new computing facility, with copper gigabit blades to support the new analysis and farm nodes and fiber uplinks to insure good connectivity to the new facility. In 2005, we'll add at least 5 additional blades. In 2006, we are likely to need to upgrade CAB to have gigabit connections compatible with the machines purchased to replace the current installation. The CAB switch would also likely need a new supervisor module.

	2003	2004	2005	2006
Analysis CPU	$505,400	$339,000	$522,000	$337,000
Primary Reconstruction	$200,000	$407,507	$423,464	$303,573
Reprocessing	NA	$1,969,574	$1,767,617	$1,477,181
Monte Carlo	NA	$446,940	$105,485	$70,323
File Servers/disk	$262,000	$360,000	$230,000	$140,000
Mass Storage	$280,000	$230,000	$100,000	$500,000
Infrastructure	$244,000	$240,000	$200,000	$220,000
FNAL Total	$1,491,400	$1,576,507	$1,475,464	$1,500,573
Virtual Center Total		$2,404,105	$1,996,482	$1,922,610

Table 6.2 Final cost estimate for FNAL and the virtual center. The FNAL spending for 2003 is also shown.

Table 6.2 shows the total estimates for 2004-2006. The FNAL total does not include the MC or reprocessing costs. Those costs are assigned to the virtual center. We assume that the MC machines are needed exclusively for DØ, but are used as part of the yearly reprocessing project. Since we assumed a 90 day duration for the re-reconstruction, we assume that those machines are used for other experiments the rest of the year, and assign 25% of the cost to the virtual center. The cost of the virtual center will evolve. Only costs which serve the entire DØ collaboration are assigned to the virtual center. In this version of the planning, the only analysis cost that is assigned to the virtual center is the FNAL resource, and other analysis resources (such as those in use at IN2P3 and GridKa Centers) are considered "user desktop" and not counted. This likely will not remain the case as DØ evolves to using grid tools and expands the use of global resources.

The spending estimates are well motivated by the use patterns that we currently observe on our systems, and the FNAL spending is in line with this year's budget. These estimates above do not reflect generous computing; in particular, with the planning shown here, DØ is not able to re-reconstruct the total data set with a single version of the reconstruction. DØ's strategy of using thumbnail datasets enables the collaboration to make good use of disk space and analysis computing.

In conclusion, DØ computing is operating well, and we are expanding the role of distributed computing within the experiment. Over the next year, we plan to replace the legacy SGI machine at FNAL, work with the CD in the context of SAM-Grid, and to establish a grid strategy that will be implemented over the course of the next 2 years.